SECURITY AND EVENT LOGGING, ANALYSIS AND MONITORING SOLUTION ARCHITECTURE ALTERNATIVE
(UTAH STATE TAX COMMISSION)

February 7, 2008
Prepared by: Robert Woolley, Chief Technologist and Strategic Planner
Reviewed by: Dave Fletcher, Kevin Van Ausdal and Michael Casey

INTRODUCTION
The objective of this document is to suggest a solution architecture for the Utah Tax Commission (UTC) that can be presented to the IRS on or before February 28, 2008. The solution must meet the logging requirements for the Tax Commission and needs to be consistent with an overall State direction for logging, analysis, and monitoring.

DELIVERABLES
The principle deliverables are the installation of key technology components for meeting the specific logging and analysis requirements established by IRS Publication 1075 and associated audit and logging requirements.

BUSINESS REQUIREMENTS
The UTC has established the following business requirements which are based upon the IRS 2006 Audit Findings, and Publication 1075:

> Audit requirements apply to operating systems, database systems, applications, and files. The audit trail should maintain a record of system activity both by system and application processes and by user activity of systems and applications. Auditable events include:
>
> - Logon/Logoff
> - Change of Password
> - Opening Files
> - Closing Files
> - All System Administrator Actions
> - All Security Administrator Actions
> - Switching Accounts or Running Privileged Actions from Another Account
> - Creation/Modification of Super-User Groups
> - Clearing of the Audit Log File
> - Startup/Shutdown of Audit Functions
> - Use of Identification/Authentication Mechanisms

- Deletion of Objects from a User's Address Space where Temporary Files have been Created (applicable to Unisys and IBM mainframe systems)
- Change Online or User Permissions or Privileges
- All Dial Up Access to the System

For each auditable event, the system must record the:

- date and time of the event;
- unique identifier of the user;
- type of event;
- subject of event and the action taken;
- origin of the request for identification and authentication events;
- name of the object introduced, accessed, or deleted from a user's address space;
- role of the user when creating the event; and,
- success or failure of the event.

Application level audit trails are those that monitor or log user activities, including data files opened and closed and specific actions taken, such as reading, editing, and deleting records or fields, and printing reports.

Audit Findings

Specific Findings of the IRS audit with detailed recommendations included the following high level findings:

Finding H24: Auditing is not configured on the system.
Risk: There is not accountability on the system.

Finding H25: Syslog records are not capturing successful/unsuccessful logon events.
Risk: Unauthorized access to the system and its resources may be undetected.

Finding H26: The use of xinetd services is not logged.
Risk: Unauthorized access to the system could go undetected.

Finding H27: The system is not configured to protect audit logs from modification, tampering, unauthorized access, overwriting, or destruction.
Risk: User activities are not being monitored, so users are not accountable for their actions on the system.

Finding H31: Remote logging not enabled.
Risk: Without logging of access list violations, routine monitoring of hostile traffic coming in to the network is difficult.

Finding H33: A “deny” ACL was found without logging enabled.
Risk: Denied connections to or through the router are not being logged.

Finding H44: Remote logging is not enabled.
Risk: Logging to a remote system is not enabled. Without logging of access list violations, routine monitoring of hostile traffic coming into the network is difficult.

Audit finding and risk and discussion comments are taken directly from the audit without modification. Detailed recommendations are not reproduced in this document for security purposes, but are implicit in the previously stated business requirements.

From a functional perspective, the UTC needs to be able to collect logs, aggregate, analyze, and monitor the data in such a manner that the logging requirements are fully satisfied.

Additional Guidelines

IRS Publication 1075[1] provides the following additional guidelines for System Auditing Guidance.:

01 The audit trail shall capture all successful login and logoff attempts.
02 The audit trail shall capture all unsuccessful login and authorization attempts.
03 The audit trail shall capture all identification and authentication attempts.
04 The audit trail shall capture all actions, connections and requests performed by privileged users (a user who, by virtue of function, and/or seniority, has been allocated powers within the computer system, which are significantly greater than those available to the majority of users. Such persons will include, for example, the system administrator(s) and network administrator(s) who are responsible for keeping the system available and may need powers to create new user profiles as well as add to or amend the powers and access rights of existing users).
05 The audit trail shall capture all actions, connections and requests performed by privileged functions.
06 The audit trail shall capture all changes to logical access control authorities (e.g., rights, permissions).
07 The audit trail shall capture all system changes with the potential to compromise the integrity of audit policy configurations, security policy configurations and audit record generation services.
08 The audit trail shall capture the creation, modification and deletion of objects including files, directories and user accounts.
09 The audit trail shall capture the creation, modification and deletion of user accounts and group accounts.
10 The audit trail shall capture the creation, modification and deletion of user account and group account privileges.

[1] Publication 1075: Tax Information Security Guidelines for Federal, State and Local Agencies and Entities, US Internal Revenue Service, February 2007, p. 72-73.

11 The audit trail shall capture: i) the date of the system event; ii) the time of the system event; iii) the type of system event initiated; and iv) the user account, system account, service or process responsible for initiating the system event.
12 The audit trail shall capture system startup and shutdown functions.
13 The audit trail shall capture modifications to administrator account(s) and administrator group account(s) including: i) escalation of user account privileges commensurate with administrator-equivalent account(s); and ii) adding or deleting users from the administrator group account(s).
14 The audit trail shall capture the enabling or disabling of audit report generation services.
15 The audit trail shall capture command line changes, batch file changes and queries made to the system (e.g., operating system, application, database).
16 The audit trail shall be protected from unauthorized access, use, deletion or modification.
17 The audit trail shall be restricted to personnel routinely responsible for performing security audit functions.

These guidelines and the previously discussed business requirements are reflective of logging and audit best practices and are generally applicable to all known agency logging requirements. With the exception some specific requirements for Federal Tax Information (FTI) that are included in the detailed recommendations, they represent a strong statement of logging and audit requirements.

TECHNICAL REQUIREMENTS

The following technical requirements have been identified, based upon information from UTC staff and CTO and DTS Security personnel:

General Compliance and Audit Requirements

The compliance, audit, and logging architecture must:

- Comply with log retention, review, and reporting requirements for PCI, HIPAA, FISMA or other specialized agency regulations such as IRS Publication 1075, and must facilitate compliance with varying regulatory requirements.

- Automatically identify and provide alerts on important audit events.

- Identify the following types of audit activity:

    - Changes to a system configuration.
    - Repeated authentication failures from the same login.
    - Repeated access control violations from the same login.
    - Repeated access control violations from the same host.
    - Account management activity.

- Provide analysis tools to assist and speed up the review process.

- Provide automation of audit/logging data collection, aggregation, analysis, and monitoring with minimal to no human intervention.

- Provide opportunities for integration with established Network Monitoring Services (NMS).

Data Collection and Management

Data Collection and Management must meet the following basic requirements:

- Ability to collect any type of log data regardless of source.
- Ability to collect log data with or without installing an agent on the logging device.
- Ability to "normalize" any type of log data for more effective reporting and analysis.
- Ability to "scale-down" for small deployments and "scale-up" for extremely large deployments.
- Provide the ability to control access to, and modification of, the log data, and log any access to the log data files.
- Provide an open architecture allowing direct and secure access to log data via third-party analysis and reporting tools.
- Provide a role based security model providing user accountability and access control. Integration with existing UMD/SiteMinder functions is highly desirable.
- Data collection must support cross platform log collection from multiple technologies such as routers, firewalls, switches, file servers, etc., using agent and agent-less techniques.

Log Analysis and Event Management Reporting and Alerts

Log Analysis and Event Management must meet the following basic requirements:

- Log Identification by log name and log source with normalization of log data for reporting purposes.

- Event forwarding for identified log entries having the most immediate operational relevance, such as security events, audit failures, warnings, and errors. Only the most important log entries should be forwarded as events.

- Risk-based prioritization and impact should be prioritized based on the event’s impact to operations.

- Role-based alerting should be easily configured to send alerts on critical events, or combinations of events, so the right alerts automatically go to the right individuals.

- Personalized analysis dashboards should allow users to quickly understand what is going on with drill down as appropriate. Users should be able to see and analyze the information most relevant to them and their role.

- Flexible reporting should be provided so users can easily supplement standard reports with custom reports tailored to defined analysis and reporting needs.

Security Integration

Logging and auditing functionality should integrate with the following security requirements:

- Central Security Monitoring
- Intrusion Detection
- File Integrity Monitoring
- Intrusion Corroboration
- Alarming and Notification

Hosting

- The application may be either server or appliance based, but must be easily scalable to support large numbers of events.
- The application must not impose undue performance loads on monitored servers and devices.
- The application must have sufficient network or be directly connected storage to meet all database, reporting, and logging storage requirements.
- The application must integrate with existing State access control mechanisms and meet all specified State security requirements for sensitive and critical infrastructure.

ARCHITECTURE

Log management[2] infrastructure typically includes the following three architectural tiers:

- **Log Generation**—Tier 1 includes the hosts that generate log data. Some servers generate log data and use the network to move that data to logging servers. Other systems make data available through other means. Log data is stored using Tier 2 resources and infrastructure.

- **Log Analysis and Storage**—Tier 2 includes log servers that receive log data from Tier 1. Data is typically transferred to the logging server in real time or using

[2] Kent, Karen, and Murugiah Souppaya, *Guide to Computer Security Log Management Recommendations of the National Institute of Standards and Technology*, NIST Special Publication 800-92, September 2006, p. 27-28.

some defined batch methodology. Log data can be stored on log or database servers.

- **Log Monitoring**—Tier 3 contains consoles that may be used to monitor and review log data from Tier 2 and provide some level of automated analysis. Log monitoring consoles are utilized to generate reports.

A simplified high level view of the logging architecture is illustrated in Figure 1. This approach incorporates either an agency centric solution when key server infrastructure is agency based, or a data center solution for hosted environments that are managed at the Salt Lake City and Richfield data centers.

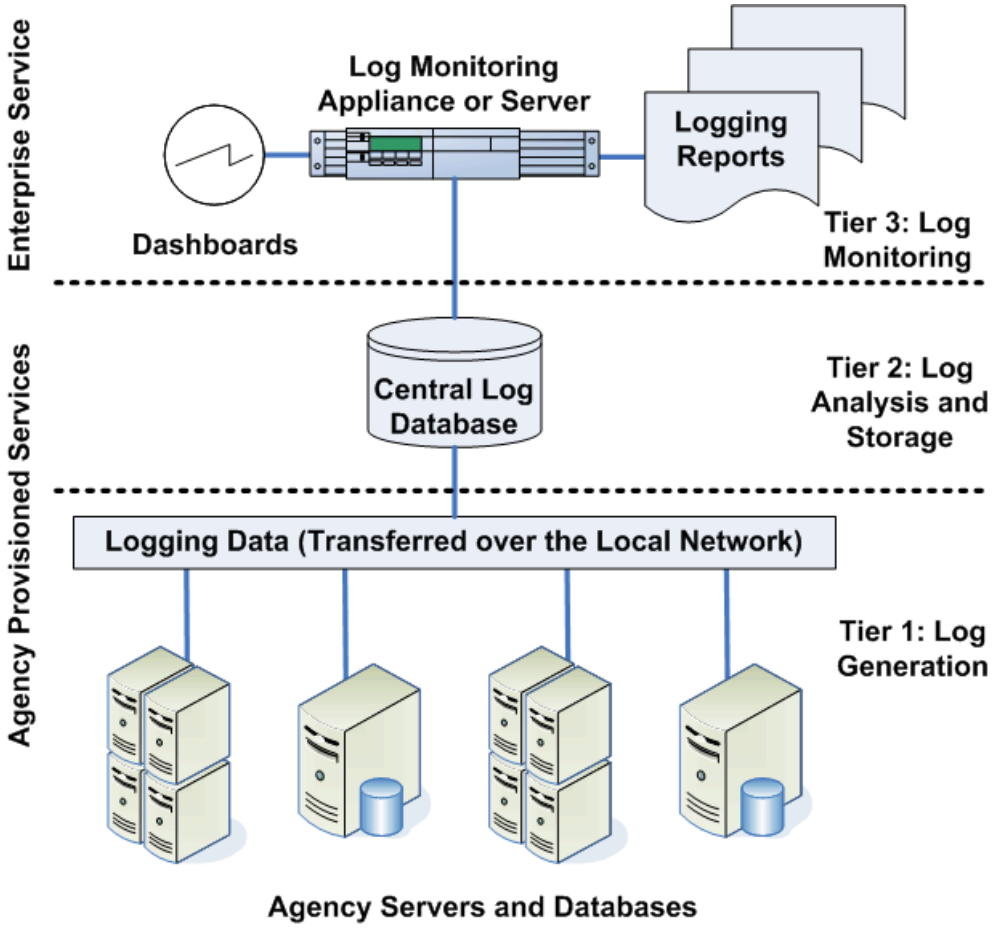

Figure 1. Three Tier Logging Architecture

Log Analysis and Storage

The second tier, log analysis and storage, can vary greatly in complexity and structure. The simplest arrangement is a single log server that handles all log analysis and storage functions. More complex second tier arrangements include:

- Multiple log servers that each perform a specialized function, such as one server performing log collection, analysis, and short-term log storage, and another server performing long-term storage.

- Multiple log servers that each perform analysis and/or storage for certain log generators. This can also provide some redundancy. A log generator can switch to a backup log server if its primary log server becomes unavailable. Also, log servers can be configured to share log data with each other, which also supports redundancy.

- Two levels of log servers, with the first level of distributed log servers receiving logs from the log generators and forwarding some or all of the log data they receive to a second level of more centralized log servers. Additional tiers can be added to this architecture to make it even more flexible, scalable, and redundant.

Log management infrastructure is typically viewed in the context of computer security log data; agencies can use the same infrastructures for other types of log data. The State could choose to have a single log management infrastructure used throughout the enterprise. In most cases, a single infrastructure is not feasible for any of several reasons, including limitations on the scalability of a single infrastructure, logging occurring on logically or physically separate networks, or concern about robustness. A single infrastructure means that a failure of that infrastructure affects logging throughout the State and interoperability issues among log generators and infrastructure components.

Communications between log management infrastructure components would typically occur over the State's regular network because the hosts generating log data may be located throughout the State. However, a physically or logically separate logging network can be used, particularly for getting logs from key devices (e.g., firewalls and network intrusion detection systems that often transfer large amounts of log data) and for transferring log data between log servers.

The State could have a single log management infrastructure for the entire enterprise, but it is more common to have multiple separate infrastructures that do not necessarily interoperate. Having a single log management infrastructure can provide a single point for reviewing all of the State's pertinent log data, but for a large organization like the State, the size of such an infrastructure and the volume of data it would have to process and store typically make it infeasible.

It is more reasonable to consider multiple log management infrastructures. The scope of each infrastructure can be dictated by many factors, including the organization's internal structure, system types (e.g., a separate infrastructure for enterprise security systems), log types (e.g., a separate infrastructure for application audit logs), and facility locations.

Log Monitoring
While Tier 2 does not lend itself as well to centralization, the monitoring and reporting component does have potential for central sharing and administration, and could be shared among agencies. This tier is critical for propagating and sharing event alerts and for standard report generation. From a UTC perspective this is the architecture tier that will be the focus of the balance of this solution architecture.

UTC Technology Architecture Platform Stack

- **Use Assumptions:** UTC estimates an average transaction load capacity requirement of 500 transactions per second.

- **Application Server/Appliance:** UTC has recommended the LogRythm base appliance. The specification for the appliance suggests the following capacity:

    - Events per second—1000
    - Raw On-line Log Capacity—100 million Events
    - Archive Online Log Capacity—1 billion Events

- **Deployment Environment:** The appliance could be deployed at the UTC premise or in the Salt Lake City data center.

- **Scalability:** Horizontal scaling by adding additional appliances is available. Specific appliance offerings from LogRythm include the following for larger volumes:

    - Midrange Appliance
    - Events Per Second—3000
    - Raw On-line Log Capacity—300 Million Events
    - Archive Online Log Capacity—2 Billion Events
    - High End Appliance
    - Events Per Second—7500
    - Raw On-line Log Capacity—600 Million
    - Archive Online Log Capacity—10 Billion

- **Software:** No special software requirement exists since this is an appliance based solution.

- **Database Repository:** To be determined, but most implementations use a MySQL back end with low cost storage. The database is used to normalize log data for analysis.

- **Data Storage Platform:** UTC SAN storage or other network attached SAN and/or tape storage at the Salt Lake City data center. Storage needs to be low cost and only a minimal amount of data is needed for current analysis.

- **Reusable Web Services:** There are no specific Web service requirements.

Engineering Design Requirements and Validation

The LogRythm solution appears to meet UTC requirements and provide substantial room for growth. There are open questions concerning integration with the older DET Mars environment. Additional work is needed to look at integration issues and the potential for LogRythm to function as an enterprise resource for analysis and monitoring.

Solution Architecture Design

The solution architecture design has been incorporated into Figure 1. A detailed design is needed to represent integration with other types of relevant data such as the MIBs from network devices. Based upon an overall assessment of deployment requirements, there do not appear to be any significant engineering or operational issues for implementing an appliance based solution such as LogRythm.

Alternative analysis

LogRythm as recommended by the UTC is a capable solution and will meet the specific requirements for the UTC with minimal issues. Reporting and scaling also appear to be adequate for combinations of agencies with similar logging requirements (e.g., Workforce Services, Health, Human Services, etc.). The appliance also looks as if it will scale well beyond the current capacity of the existing MARS installation by DET.

Other strong solution alternatives are available and should be considered. A short list of other alternatives that seem most applicable to the State would include:

- eIQ (See http://www.eiqnetworks.com)—This solution is currently being reviewed and tested by the Department of Alcohol and Beverage Control (DABC) and represents a very comprehensive SIEM monitoring and reporting solution that will also scale to meet Tax and enterprise requirements.

- MARS Upgrade—The existing MARS environment will not scale much beyond where it is today, and is somewhat dated. None the less, the MARS offering from Cisco is a useful analysis tool that could also meet some, if not all, of the UTC requirements. No cost data has been identified for upgrading this environment.

- Novell Sentinel—Sentinel 6 has received favorable reviews from Gartner and a number of other reviewers in the SIEM marketplace. The product is server based and provides a robust set of logging tools and reports. The product is available under the existing master license agreement with Novell. A single server instance is $32,500 with annual maintenance at $6,650. Additional server licenses are available at $8,750.00 per server plus $2,190 in annual maintenance. Server deployment uses standard SUSE Linux servers.

An appliance based solution would appear to be in the best interest of the State from an ongoing maintenance and operations perspective. With the exception of MARS, which is available under the Cisco contract, all other solutions raise some procurement concerns. The UTC has resources available for LogRythm, but may not have sufficient resources for an eIQ solution.

Other open issues include the need to integrate the SIEM solution with existing, although somewhat dated, Network Monitoring Services (NMS). A conclusion as to what the logging product should incorporate on an enterprise level, and upon what cost basis would be required for an enterprise service, whereas an agency point solution is much less complex.

Preliminary Cost Analysis

This cost analysis is incomplete since costs for eIQ and Mars alternatives have not been determined reliably. All cost figures are based on estimates from LogRythm and would likely be less in a competitive procurement environment. Personnel costs are excluded from the cost analysis.

Year One

Assumption—Tax only at an entry level configuration

Base Appliance		$15,000
Events per second—1000		
Raw On-line Log Capacity—100 million		
Archive Online Log Capacity—1 billion		
Storage (Estimate based on UTC only)		$20,000
	TOTAL	$35,000

Ongoing Expenses

Base Appliance Maintenance (Estimated)		$ 3,000
Storage (Estimate based on UTC only)		$25,000
	TOTAL	$28,000

A UTC only solution could be procured within existing funding, and could be procured with a competitive bid process. No RFP would be required. From a risk perspective, the solution may have to be scaled up quickly since capacity estimates are based on a best guess for event volume per second. If that were the case, the appliance cost would double from $15,000 to $30,000. The other configuration options detailed in the platform stack discussion are $30,000 and $45,000 respectively.

An enterprise approach would require the release of a competitive RFP, which will require additional time and identification of budget resources prior to the RFP release. Some limited additional requirements gathering would be needed, but for the most part,

the Tax requirements are inclusive of most other common agency logging needs. The UTC needs to identify and implement a solution as quickly as possible, and before the next audit cycle with the IRS. This tends to favor the LogRythm recommendation.

At a general level, the top three alternatives reflect the following costs:

Product	Appliance/Server	License Cost	Maintenance	Total
LogRythm	$15,000		$ 3,000	$18,000
Sentinel	$ 7,500	$32,500	$ 6,650	$46,650
eIQ	$43,195	$36,326	$15,903	$81,424

All of the logging options require similar storage and connectivity, so those costs have not been included in the above numbers. Personnel costs are also considered to be somewhat equivalent and have not been detailed. eIQ costs are based on a quote to DABC for a similarly sized environment. The options have been compared in terms of the UTC implementation only. Of these three options LogRythm appears to be the least cost solution, while Sentinel is the simplest contract and procurement alternative.

Summary and Recommendation

Centralized logging services, including monitoring and analysis, are a prime area for providing value added enterprise services to agencies. DTS has an opportunity to provide analysis and monitoring services for all agencies. There are substantial gaps in the number of servers and applications that are currently monitored, and existing logging and analysis infrastructure is becoming dated. Time to benefit is an important consideration for the UTC solution. DTS must determine if this will be a point solution with future enterprise service implications for other agencies. The LogRythm solution fits into the described logging architecture and does not represent any architectural issues, either as a point solution or as a component of a multi-agency or enterprise monitoring approach. DTS needs to make this determination and then address the following issues suggested in the TA Review for Logging.

- Develop a product definition and scope for logging services.
- Define the funding mechanisms for supporting centralized logging services at the least cost to agencies.
- Assess the existing server population to identify what servers should be added to centralized logging pools.
- Meet with agencies and document specific logging requirements, including those required by law or rule from other entities such as the federal government.
- Establish processes to integrate alerts and events from agencies that are doing centralized logging so they can be integrated with other enterprise alerts and events.
- Assess the cost, performance, and capabilities of LogRythm, eIQ, and Sentinel in terms of meeting State requirements.

- Update the existing DET central logging environment with the objective of providing additional services to agencies that have specific logging requirements.
- Determine if the DET central logging and monitoring environment could be implemented more cost effectively with LogRythm or one of the other alternatives as opposed to upgrading the MARS environment.
- Release a competitive RFP, if needed, for the procurement of SIEM software or appliances for enterprise logging services analysis and monitoring and procure needed infrastructure.

Centralized logging services require additional investment to be effective. Failure to invest on an enterprise level pushes added expense to agencies, which should be avoided. None-the-less, some aspects of a comprehensive logging solution on tiers one and two of the logging architecture may be better handled on an agency level to minimize network traffic.

An appliance based SIEM tool seems to represent a least cost point of entry for DTS that will meet agency needs and add value. The UTC solution does not preclude integration with a future enterprise product for log monitoring and reporting. The technology utilized is on an enterprise level and will be dependent upon RFP results and the capacity and desire of agencies to pay for log monitoring. It may not be reasonable to delay UTC based on a logging product for DTS that has not been defined and may not be financially viable.